AAS 93-603

HIGH-PRECISION MAGELLAN ORBIT DETERMINATION FOR STEREO IMAGE PROCESSING

P.W. Chodas*
S.A. Lewicki†
S. Hensley††
W.C. Masters'

The Magellan spacecraft has systematically imaged the surface of Venus using Synthetic Aperture Radar (SAR) for two years. During this time Venus rotated three times under the orbit, causing the ground track to make three cycles across the surface. The radar was operated in such a way that terrain seen in the first cycle was imaged again in the third cycle at a 10-20° smaller incidence angle. Because of the different viewing angle, overlapping images from the two cycles can be combined to produce stereo images and high-resolution digital elevation maps. Stereo processing is very sensitive to ephemeris errors, however, and the current Harlh-based ephemeris produces large artifacts in stereo products. This paper describes a technique for improving the Magellan ephemeris on multiple blocks of orbits using measurements of landmarks. The technique is particularly appropriate for generating ephemerides for usc in stereo processing because it reduces relative ephemeris errors between non-contiguous data arcs. When applied in a test of stereo processing methods, the technique produced ephemerides with relative errors of about 100 In, an order of magnitude smaller than those in the best Ear[h-based ephemerides.

INTROD1JCTION

From September 1990 to September 1992, the Magellan spacecraft systematically observed the surface of Venus using a Synthetic Aperture Radar (SAR). During that time the spacecraft's ground track swept across Venus three times, as the planet completed three 243-day rotations. Each sweep of the ground track was referred to as a "cycle" of the mission. To maximize science return, the radar was operated differently on each cycle. During the first and third cycles, the surface was viewed in a left-looking mode, i.e., the radar was pointed left of the ground track as seen by an observer facing in the direction of the motion. The two cycles differed, however, in the off-nadir angle, referred to as the

* Member Technical Staff, Navigation Systems Section, Jet Propulsion Laboratory, California Inst itute of Technology, Pasadena, California 91109.
† Member Technical Staff, Earth Observat ion and Analysis Systems Section, former'ly of the Image Processing Applications and Development Section, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California 91109.
†† Member Technical Staff, Radar Science and Engineering Section, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California 91109.

"look angle", which varied with altitude in both cycles. The look angles in cycle 1 were as large as possible, to maximize resolution. Smaller look angles were used in cycle 3 to provide a view of the same terrain at 10-20° smaller incidence angles, so that stereo images could be produced by combining overlapping cycles 1 and 3 images. Although spacecraft transponder problems limited the cycle 3 coverage, a large archive of images with the smaller incidence angle was acquired, covering about 25°/0 of the planet's surface.

The processing of radar echo data into images requires precise knowledge of the spacecraft's orbital ephemeris. Magellan's orbit is determined to sufficient precision using Earth-based Doppler tracking measurements. However, when radar images from individual orbits arc combined into mosaics necessary for geological mapping, relative errors between orbit solutions arc sometimes evident as discontinuities running through the images. Stereo processing is particularly sensitive to ephemeris errors: even small relative ephemeris errors between two orbits observing the same terrain cause relatively large artifacts in stereo products such as Digital Elevation Maps (DEMs). The standard ephemeris determined from Earth-based data is generally not precise enough for accurate stereo processing.

Reference 1 describes a technique which improves the precision of Magellan orbit determination by incorporating measurements of landmarks visible in the radar images. Being Venus-relative, the landmark measurements provide information complementary to that supplied by Earth-based tracking measurements. This paper describes how the ephemeris improvement technique has been extended to handle multiple orbit blocks linked by common landmarks. The new technique is particularly appropriate for stereo processing because the ephemeris can be improved on blocks of orbits from both cycles 1 and 3 in a way that reduces the relative intercycle ephemeris errors. "l'he technique was used successfully to improve cycle 1 and cycle 3 ephemerides over a test region in order to try various stereo processing methods, and the resultant DEMs contained no apparent artifacts due to ephemeris errors.

THE MAGELLAN MISSION

The primary goal of the Magellan mission was to perform high-resolution radar mapping of the surface of Venus on a global scale. The original goal of achieving a 70°/0 coverage of the surface was surpassed in the first 8 months of radar mapping, and over the course of two years of radar operations, over 98°/0 of the surface was mapped. Magellan's orbit about Venus is near-polar and, until recently, moderately eccentric, with a periapsis altitude of about 290 km and apoapsis altitude of about 8460 km. (Although Magellan's orbit was lowered and made near-circular through the use of aerobraking during the period May-August 1993, the work discussed in this paper applies to the original orbit.) The radar operated for only a 37-minute period around the periapsis of each cycle 1 orbit, during which time the high-gain antenna was pointed towards Venus and the radar echo data was recorded on board. During the remainder of each 3.26-hour orbit, the high-gain antenna was pointed at the Earth and the radar data was played back at 1/3 the speed it was recorded.

The radar data acquired on each orbit covers a long narrow North-South strip of surface. These swaths are only about 23 km wide, but up to 15,000 km long. Planetary rotation from onc orbit to the next causes each swath to lie just to the cast of the previous one, with a small amount of overlap, since Venus rotates 21.25 km per Magellan orbit, at the equator, and less at higher latitudes. During cycle 1, the latitudinal coverage of mapping was maximized by alternating the latitude coverage of the swaths: on even-numbered orbits the swaths ran from the north pole to about 52° South, while on odd-numbered orbits they started at about 54° North and ran to 78° South, This strategy produced no gaps in coverage because at high latitudes the linear displacement of the swaths duc to planetary rotation was small enough that even alternate swaths overlapped. Due to a transponder problem which lowered the downlink data rate in cycle 3, the swaths were half the length of cycle 1 swaths, and the alternating latitude coverage scheme was not used.

THE MAGELLAN SYNTHETIC APERTURE RADAR SYSTEM

Radar mapping is performed by flying an antenna with a very narrow beamwidth, pointing it to onc side of the nadir in the plane perpendicular to the horizontal component of the velocity, transmitting very short pulses of radio energy, and sampling the echo of each pulse returned from the ground. The echo is spread out in time because the various ground points illuminated by the beam are at varying distances from the antenna; the echo is spread out in frequency due to Doppler shifting, and the various ground points have varying radial velocities, as seen by the radar (for example, points ahead of the spacecraft have positive radial velocities while points behind the spacecraft have negative radial velocities). The antenna is pointed to one side of the nadir to avoid the ambiguity of which side of the ground track the echoes come from.

High resolution radar mapping requires very narrow antenna beamwidths. Since the beamwidth at a given wavelength is inversely proportional to the antenna aperture, it would seem that high-resolution mapping requires impracticably large antennas (hundreds of meters in size). Synthetic Aperture Radar provides a solution to this problem by synthesizing a large aperture antenna from the forward motion of a much smaller real antenna. To achieve this, the radar transmits a coherent series of several hundred pulses, called a burst, and sums the echo returns as if received by a single antenna with an aperture equal to the distance the radar traveled during the series of pulses.

The radar takes several hundred samples of the phase of the returned echo from each pulse as it is received. Each sample is stored in a "range bin" characterized by the time delay from the pulse transmission to the reception of the sample, which determines the location of a strip of ground at a constant range from the antenna. Over the course of a burst, a series of phase samples is collected for each range bin (one sample for each pulse in the burst). The phases of the echo samples are adjusted to account for the position offsets between the real antenna at the pulse times and the position of the synthesized antenna. Once the echoes from all the pulses in a burst have been collected, a frequency analysis is performed via a Fast Fourier Transform (FFT) on the series of phase samples

for a given range bin, The amplitudes of the transformed signal at various frequencies represents the radar reflectivity at various ground points within the strip, since the echo from each point has a unique Doppler shift determined by its position ahead of or behind the synthesized antenna. By repeating the FFT for each of the range bins, a radar reflectivity image in range/Ilopplcr space is built up. This raw range/Doppler image is called a "look"; each burst produces one look.

By themselves, individual looks arc too noisy to be of much USC. Since SAR is a coherent imaging technique, the looks contain a large amount of speckle noise, a result of random interference within each range-I)opplcr resolution element. The speckle noise can be greatly reduced by taking several looks of each patch of surface, registering the looks to a common frame and averaging pixels. in particular, the individual looks arc designed to overlap considerably in the along-track direction: each surface point in a Magellan swath is typically seen in at least 4 looks, and as many as 18. The frame in which the looks arc merged is a latitude/longitude frame placed on a reference surface defined by a low-frequency topographic model determined from earlier missions to Venus. The pixel size of the projected image is 75 m, about half the actual radar resolution, which is 100-300 m in the range direction and 120 m in the along-track direction (i.e., the pixels arc oversampled after projection).

During a mapping pass, the Magellan SAR transmitted about 5500 bursts of radar, averaging over 2 bursts per second. When merged in the latitude/longitude frame, the 5500 individual looks become a single long North-South image about 300 pixels wide by 200,000 pixels long. This is the basic image produced on each orbit. Because of their size, the basic image swaths arc unwieldy to work with. More useful images arc produced by mosaicking series of swaths at selected latitudes, The mosaics subtend an area of surface about 5 degrees on a side, and arc comprised of dozens of basic image swaths.

The production of radar images requires knowledge of the spacecraft ephemeris, which is needed when each look is projected into the latitude/longitude frame, Errors in the ephemeris cause the looks to be projected to incorrect positions in latitude/longitude. Although the frequency of the ephemeris errors is not high enough to cause any significant misregistration of features among the looks being merged at any given point in the swath, the ephemeris errors cause the basic images to be shifted, rotated, and/or warped in the latitude/longitude frame. Spacecraft attitude errors, on the other hand, have essentially no effect on the range/Doppler coordinates of a feature (provided the feature is still within the main lobe of the beam). Compare this with optical imaging in which the line/pixel coordinates of a feature are usually very sensitive to the attitude (pointing) of the camera.

STEREO PROCESSING OF SAR IMAGES

Stereo processing, whether in the optical or the radar regime, uses two images of the same region of surface to extract three-dimensional information about that region[2]. Although the individual images are projections of the three-dimensional scene into two dimensions, information on the third dimension can be derived from differences in the

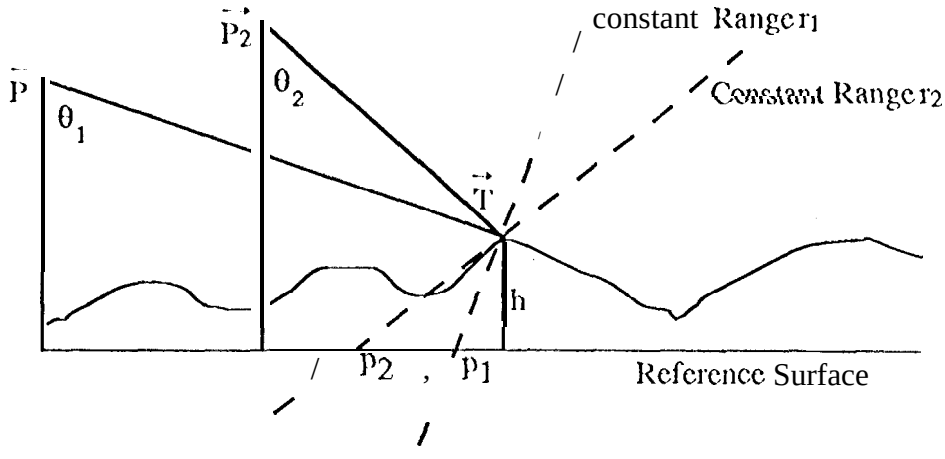

Figure 1. Simplified Radar Stereo Geometry

relative positions of features in the two images, assuming the viewing directions of the two images arc different. To illustrate how this works for radar images, consider the simplified geometry illustrated in Figure 1. The plane of the figure is perpendicular to the reference surface. Assume the spacecraft is at position P_1, relative to the center of Venus, at the time of the first image, and, for simplicity, assume the spacecraft is also in the plane of the figure at the time of the second image, at position P_2. Because radar is side-looking, the spacecraft motion must be nearly perpendicular to the figure. Consider a surface feature at position T. The radar measures the ranges (denoted by r_1 and r_2) from the spacecraft to the feature at the two times. Assuming these ranges are large, the loci of constant ranges from the spacecraft can then be viewed locally as straight lines perpendicular to the line-of-sight vectors, as illustrated by the dashed lines in Figure 1. In the production of the radar images, all features are projected onto the reference surface. The projections of the feature at T in the two images are the intersections of the constant range lines with the reference surface, denoted by p_1 and p_2 in the figure. The difference in the feature location in the two images is called the parallax, and from it the height of the feature above the reference surface can be derived using simple trigonometry,

$$h = \frac{p_2 - p_1}{\cot \theta_2 - \cot \theta_1} \qquad (1)$$

where θ_1 and θ_2 arc the angles at the spacecraft from the local vertical to the line-of-sight vectors, as shown in the figure.

The method described above has used several simplifying assumptions, including that P_2 lies in the same plane as P_1 and T. As a result, the Doppler coordinate of the feature was ignored. A more general height dct ermination, and indeed a det erminat ion of the full 3-dimensional position of the feature, can be obtained by using the original range/I) oppler

coordinates of the feature instead of latitude/longitude. This is possible because the polynomial coefficients originally used to map each look from range/Doppler to latitude/longitude arc stored as ancillary data to the basic image swath, The observed ranges and Dopplers can be related to the positions and velocities of the spacecraft at the two image times, and the position of the feature. let **b** be the position of the surface feature relative to the center of Venus, measured in the Venus body-fixed frame. Consider only a single observation for now, and let **r** and v be the inertial-frame position and velocit y of the spacecraft relative to the center of Venus at the observation time. With the effects of atmospheric refraction ignored for simplicity, the observed range p and Doppler f_D are given by

$$\rho = |\mathbf{C}_b\mathbf{b} - \mathbf{r}| \tag{2}$$

$$f_D = \frac{-1}{\lambda\rho}(\mathbf{C}_b\mathbf{b} - \mathbf{r})^T(\boldsymbol{\omega}_V \times \mathbf{C}_b - \mathbf{v}) \tag{3}$$

where C_b is the rotation matrix from the inertial to the body frame at the time of the observation, λ is the radar wavelength, and $\boldsymbol{\omega}_V$ is the angular velocity of Venus. Repeating this pair of equations for each of the two stereo observations, one obtains an overdetermined system of four nonlinear equat ions in the three unknowns **b**, referred to as the radar stereo equations. The system is solved using nonlinear least squares techniques.

Regardless of which method is used to derive stereo height, the problem of matching features from one image to the other at the subpixel level remains. This is accomplished using image correlation techniques. For a given small region, usually square, on one image, a box of the same size is moved around on the second image until the contained images are maximally correlated. SAR image matching differs from optical image matching in several ways. The SAR images contain speckle noise, even though it is greatly reduced by averaging pixels from multiple looks. The noise obscures fine details in the scene texture, which aid the correlation process for optical images, and makes correlation of scenes without features impossible. Another problem is the geometric distortion which arises when the same region is imaged from different viewing angles. These distortions can be so severe that a feature clearly visible in one image is impossible to identify in another image with a different view angle.

Various algorithms have been developed for performing image matching. In hierarchical matching techniques, the two images are repeatedly correlated with uniformly smaller and smaller correlation boxes, the pixel shifts at one level used to initialize the correlation search at the next lower level. 1 n multi-resolut ion pyramidal approaches[3], the size of the correlation box size stays fixed while the scale of the image changes from one level to another by combining pixels; large features are still matched before the detailed terrain. Whatever the approach, the output of the scene matching is called a disparity map, which gives, for each stereo resolution element in one image, the two-dimensional offset to the matching clement in the second image. For Magellan data, the stereo resolution elements can get as small as about 300 m, or 4 image pixels, in size, Once the

disparity map is available, the heights of all the resolution elements arc determined using one of the algorithms described above, and a digital elevation map is produced.

BASELINE MAGELLAN ORBIT DETERMINATION

The baseline technique for determining Magellan's orbit used Earth-based Doppler measurements of the spacecraft velocity acquired during the periods when the spacecraft ant enna was point cd at the Earth[4]. Two t ypcs of Doppler measurements were used: two-way Doppler, which is sensitive to the spacecraft velocity along the line of sight, and differenced Doppler, formed by differencing simultaneous Doppler measurements from two stations, which is sensitive to spacecraft velocity perpendicular to the line of sight. Measurements of' the second type were available only occasionally. The spacecraft position can be determined using these measurements 10 an absolute accuracy, relative to the center of Venus, generally better than 10 km in the horizontal plane and 300 m radially, The relative position accuracy (one orbit with respect to the previous orbit) is considerably better: usually less than 1 km horizontally and 150 m radially. Although this level of accuracy is adequate for processing the radar data, it is not sufficient to eliminate artifacts in image mosaics and stereo products. For example, a 1-km along-track relative error between consecutive orbits causes a noticeable discontinuity in an image mosaic because the resolution of the image is an order of magnitude better.

Several factors caused the orbit determination accuracy to deteriorate at various times. During periods of poor viewing geometry, the observability of Magellan's orbit using predominantly line-of-sight observat ions was limit cd. For example, when the Earth-spacecraft line of sight was nearly in the orbit plane (i.e., the orbit was seen "edge on"), the orbit inclination was poorly estimated. Similarly, when the orbit was seen face-on, the in-plane elements were poorly determined. When the latter occurred in March 1991, radial ephemeris errors increased threefold to about 0.4 km.

The accuracy and availability of the Doppler measurements also varied, The tracking measurements were not always made at X-band, the preferred frequency. When S-band was used for either uplink or downlink, the measurement accuracy decreased by a factor of 5. Furthermore, near superior conjunction, when Venus was on the opposite side of the Sun from the Earth, the Doppler measurements were corrupted by the increased solar plasma along the signal path, and the measurement accuracy degraded by a factor of up to SO. Relative ephemeris errors of up to 5 km were seen during this period.

Relative orbit-to-orbit ephemeris errors arc largest across so-called “navigation boundaries”, i.e., the boundaries between the blocks of 7-8 orbits covered by each navigation solution, The spacecraft ephemeris within each block (or data “arc”) is computed via a single continuous numerical integration of the equations of motion, and is based on a single set of tracking observations. Relative ephemeris errors across navigation boundaries are larger than those within a navigation solution because the ephemerides arc computed from different numerical integrations and arc based on different sets of tracking observat ions,

Magellan mosaics are comprised of dozens of orbits, and therefore contain several navigation boundaries, Since the relative ephemeris errors across these boundaries are often significantly larger than the resolution of the image, noticeable discontinuities often occur along the boundaries. These artifacts arc particularly troublesome in mosaics of images acquired near superior conjunction, where ephemeris errors were largest. But even smaller 1-km relative errors which occur during favorable geometries are quite discernible in the mosaics.

Stereo processing is very sensitive to ephemeris errors: even small errors lead to noticeable artifacts, Figure 2 illustrates this problem using real Magellan data. It shows the disparity maps for a small region near Maxwell Montes, created by stereo-matching image mosaics from cycles 1 and 3 produced using the standard Earth-based ephemeris, The top image shows cross-track shifts: dark areas represent shifts to the right, light areas shifts to the left; the bottom image shows along-track shifts, with lighter areas indicating shifts downwards. Although the ridge features in the cross-track image arc duc to real parallax of features, the strong vertical banding is caused by orbit errors between navigation solutions. The wide vertical bands arc 6 or 8 orbits wide and correspond to navigation solutions; the edges of the bands mark the navigation boundaries. Bands from the two cycles arc superimposed: the navigation boundaries from cycle 3 arc slightly less vertical than those from cycle 1. The largest along-track discontinuity, running down the center of bottom image, is across a cycle-1 navigation boundary, and corresponds to a relative ephemeris error of about 9 pixels, or 700 m. The cross-track relative errors are somewhat smaller, but they have a dramatic effect on the stereo processing, producing artificial cliffs almost a kilometer high running down the length of the DEM. Artifacts like these cannot easily be removed cosmetically from the stereo products -- improving the accuracy of the spacecraft ephemeris is the best solution to the problem,

IMPROVED ORBIT DETERMINATION USING SAR LANDMARKS

The fact that ephemeris errors are noticeable in radar mosaics indicates that the Magellan SAR images have high enough resolution to cent ribute orbit information, Measurements of distinct features ("landmarks") in Magellan radar images provide a means for improving the accuracy of the spacecraft's ephemeris. A method for combining landmark measurements with the standard data set of ground-based Doppler measurements to compute an improved spacecraft ephemeris has been developed, and is described in Ref. 1. The technique has been demonstrated to significantly improve the accuracy of the orbit estimate. The landmarks provide Venus-relative information which helps to tie orbits together and reduce relative ephemeris errors. Onc difficulty with using landmark measurements, however, is that the Venus-fixed coordinates of the landmark (latitude, longitude, and radius from the center of Venus) are not well known, and must bc estimated along with the spacecraft orbit parameters. In order to provide orbit information, a landmark must bc observed at least twice, (each observation providing two components of information), and it must therefore lie in an overlap region between two swaths.

<——CYC. 3 NAV. SOLUTION->
<——CYC. 1 NAV. SOLUTION—-->

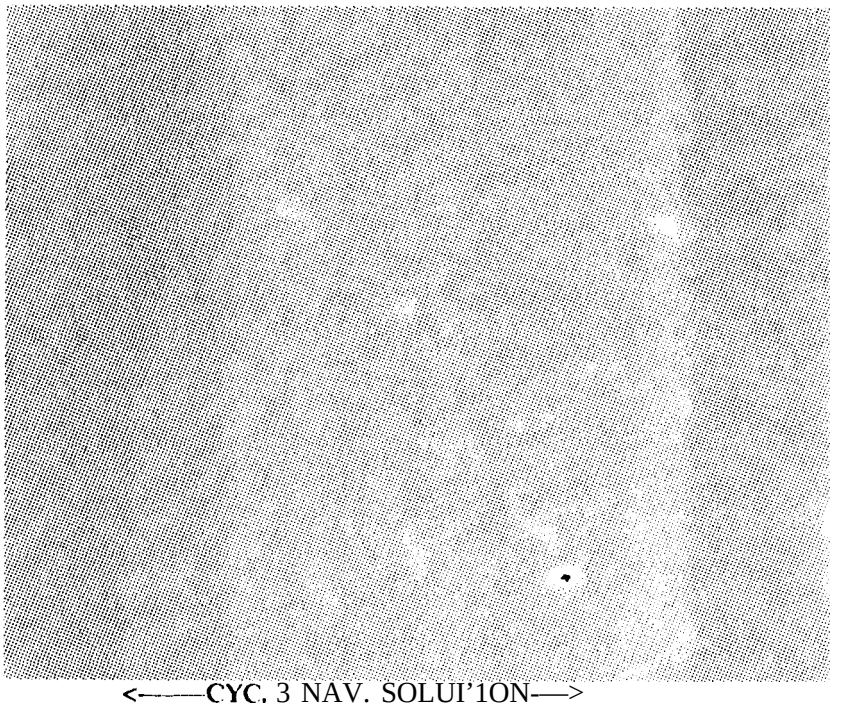

<——CYC. 3 NAV. SOLUI'1ON—>
<——CYC. 1 NAV. SOLUTION-------->

Figure 2. Disparity Maps for Maxwell: Cross-'l'rack (top) and Along-Track (bottom)

Landmarks arc selected manually via an interactive program which displays the same latitude region from two swaths, Once a suitable feature in the overlap region has been identified and a correlation box size selected, the program performs an image correlation similar to that used for stereo image matching. The program records the pixel coordinates of the center pixels of both correlation boxes and assigns an identification number to the landmark. landmarks selected on previous pairs of orbits arc called up and correlated on as many ncw orbits as possible. Landmarks north of 80° North latitude can be measured on 4 or more alternating swaths, The number of landmark measurements per orbit averaged about 10. Once a landmark's pixel coordinates arc identified on a given swath, the program reconstructs the range/Ilopplcr coordinates of the landmark using the same technique described above for the radar stereo equations. Basically, a burst containing the landmark is selected, the pixel coordinates arc mapped back to range/Doppler coordinates, and the measurement is assigned the time tag of the center time of the burst. Ref. 1 describes the processes of selecting, correlating, and processing the landmarks in more detail.

The basic steps taken in processing landmark measurements are the same as for other navigation measurements[5]. The spacecraft trajectory is numerically integrated over the entire data arc using a priori values for the initial orbit parameters and other dynamic parameters. Expected values of the measurements arc computed using expressions analogous to Eqs. (2) and (3), with a priori values for the landmark coordinates and Venus rotation model parameters. 'J'he effects of atmospheric refraction are accounted for. The partial derivatives of the expected measurement values with respect to all estimated parameters arc also computed. Simplified versions of these partials were given in Ref. 1. Residuals are formed by subtracting the expected values from the observed values. The landmark residuals and partials arc combined with the ground-based tracking residuals and partials in a linearized, weighted least squares procedure which minimizes the sum of squares of the residuals, $\mathbf{Ax} - \mathbf{z}^2$, where $\mathbf{z}$ is a vector containing all the measurement residuals, $\mathbf{A}$ is the matrix of corresponding partials, x contains the solved-for corrections to the estimated parameters, and the measurement weights have been omitted for clarity. The problem is invariably overdetermined, so that $\mathbf{A}$ has many more rows than columns. The particular estimation technique used to solve the probelm is the square root information filter[6]. Instead of forming the normal matrix for the problem, this method uses Householder transformations to reduce the problem to a minimization of the quantity $\|\mathbf{Rx} - \mathbf{z}'\|$, where $\mathbf{R}$ is an upper triangular matrix referred to as the square root information matrix; $\mathbf{R}$ and $\mathbf{z}'$ satisfy

$$T[A \quad z] = \begin{bmatrix} R & z' \\ 0 & e \end{bmatrix} \tag{3}$$

where 'T' is a product of Householder transformations and $\mathbf{e}$ is vector of transformed residuals. Since $\mathbf{R}$ is square, the solution to the least squares problem is simply

$$\mathbf{x} = \mathbf{R}^{-1}\mathbf{z}' \tag{4}$$

The corrections x are applied to the estimated parameters and the steps are repeated until the corrections become sufficiently small. Three iterations are generally sufficient to attain convergence.

The measurements are weighted in the estimation by the inverse of their assumed accuracy. The SAR range accuracy was assumed to be a constant 30 m, one sigma. Since the SAR Doppler accuracy is a function of burst duration, which varies over a mapping pass, these measurements were normalized by the burst duration. The normalized SAR Doppler accuracy was then assumed to be constant at a value equivalent to a 50 m along-track accuracy (one sigma).

A priori estimates of the coordinates of the landmarks are needed to perform the estimation. A priori estimates of latitude and longitude are immediately available from the basic images, which are aligned with the latitude/longitude frame. The a priori estimate of the radius from the center of Venus is obtained by interpolating within either the Magellan altimetry data set, if available at that location, or within the Pioneer Venus/Venera 15-16 topography mode]. If the radius value is obtained from the more-accurate Magellan altimetry, it is fixed in the solution (i.e., not estimated), to reduce the dimensionality of the problem.

Other uncertain parameters affecting the SAR measurements are the direction of the spin axis of Venus, and the planet's rotation rate, which together determine ω_V. The ability to estimate these parameters was included in the software, although the observability of the pole direction was usually very weak due to the limited range of longitudes in each solution. Generally, the pole direction was not estimated, but fixed at the value computed by the Magellan Geodesy/Cartography team at the RAND Corp. The rotation rate, on the other hand, was well-determined for intercycle solutions, discussed next,

MULTI-ARC ORBIT-DETERMINATION USING landmarks

The method described above and in Ref. 1 applied only to individual orbit-determination arcs. To support stereo processing, the technique of using landmark measurements to improve the orbit determination was extended to multiple arcs. Recall that a data arc is a block of orbits over which the ephemeris is computed via a continuous numerical integration, starting from an initial state which is adjusted to minimize the residuals of measurements made during those orbits. The ephemerides for different data arcs are based on different sets of measurements, causing discontinuities in the ephemeris across the boundaries between arcs (i.e., the navigation boundaries discussed earlier). Data arcs for Magellan were typically 8 to 12 orbits long, a size which is a compromise between opposing effects. On shorter data arcs, the ephemeris is less well-determined because it is based on fewer measurements; on longer data arcs, ephemeris errors build up over the course of the numerical integration due to inaccuracies in the dynamic models.

In multi-arc orbit determination using landmarks, the ephemeris is estimated on multiple arcs simultaneously, with measurements of landmarks used to tic the arcs together. The orbit solutions on the individual arcs arc allowed to vary independently, but the variations are constrained by measurements of common landmarks. The data arcs can bc contiguous or separated by multiples of a Venus rotation period. For contiguous arcs, the common landmarks reduce the discontinuities across arc boundaries. For intercycle solutions, the common landmarks tic an ephemeris on onc cycle to that on another, reducing the relative errors between them. This makes the multi-arc technique ideal for computing improved ephemerides for Magellan stereo processing, Intercycle orbit determination using landmarks also enables an accurate determination of the Venus rotation rate.

In the multi-arc technique, the estimated parameters for each arc arc partitioned into an arc-dependent portion and a common portion, $x = [x_a \; XC]^T$. R and z' arc similarly partitioned, and on each arc, the problem becomes onc of minimizing

$$\left\| \begin{pmatrix} R_a & R_{ac} \\ & R_c \end{pmatrix} \begin{pmatrix} x_a \\ x_c \end{pmatrix} - \begin{pmatrix} z_a \\ z_c \end{pmatrix} \right\|^2 . \tag{5}$$

As measurements arc processed for an arc, the information is packed into R and z'. At the end of the arc, the arc-dependent rows of R and z' are saved, and then initialized to zero for the next arc. Information continues to accumulate in the common-parameter rows of R and z'. At the end of the last arc, the solution for the corrections to the common parameters are obtained from

$$xc = R_c^{-1} z_c \tag{6}$$

Then, for each arc, the arc-dependent rows of R and z' are recalled and the corrections to the arc-dependent parameters arc computed from

$$X_a = R_a^{-1} z_a - R_a^{-1} R_{ac} x_c \tag{7}$$

EPHEMERIS IMPROVEMENT RESULTS FOR A STEREO TEST REGION

In order to validate and compare various stereo processing techniques, two regions of Venus seen in both cycle 1 and cycle 3 images were selected for the processing of digital elevation maps. One was a region at 66° North on the western flanks of Maxwell Montes, the highest mountain on Venus, and the other was Gula Mons at 22° North. Conveniently, the two regions were covered by a single block of 25 orbits in cycle 1, and a slightly larger block of32 orbits in cycle 3. The disparity maps shown in Figure 2 covered a portion of the selected region near Maxwell; they clearly demonstrated that ephemeris improvement was needed to avoid large artifacts in the elevation maps. Since one of the stereo processing techniques was based on image mosaics, the plan was to use the ncw

ephemeris to reprocess the SAR data to create new image swaths, which would in turn be used to make new mosaics of the regions. The stereo technique based on the radar stereo equations could use the new ephemeris directly and did not require reprocessing the SAR data.

The 57 orbits requiring ephemeris improvement were split into 7 arcs, 3 in the first cycle and 4 in the third. All arcs were 8 orbits long, except the last cycle 1 arc, which contained 9 orbits. The cycle 1 data arcs cover the period Oct. 4-7, 1990; the cycle 3 arcs occurred Jan. 30- Feb. 4, 1992. Image swaths from even-numbered cycle 1 orbits extended north to the pole; those from odd-numbered cycle 1 orbits extended as far south as 79°. Due to a transponder problem which lowered the downlink data rate, the cycle 3 orbits were much shorter, extending only from 76° to 20° North. Also, the image swaths from cycle 3 were not parallel to those from cycle 1 because a different look angle profile was used; a given cycle 3 swath crosses 25 cycle 1 swaths in this region,

The landmark measurements were initially made on each cycle separately. About 340 landmarks were selected and measured on cycle 1 orbits, and 210 on cycle 3 orbits, Then, landmarks were tied across cycles at selected latitudes, viz. 10° bands at the top and bottom of the cycle 3 swaths, and a 2° band at 54° North. Being restricted to these latitude bands reduced the number of intercycle orbit pairs that had to be examined. Initial single arc solutions uncovered large SAR Doppler residuals for 3 of the cycle 3 orbits, which were due to the use of incorrect spacecraft clock calibrations. landmark measurements on these orbits were simply deleted from the solution. The total number of landmarks that remained after removal of these orbits and other clearly erroneous points was 516, with 107 landmarks observed on both cycles. The total number of landmark measurements used was 1388, and the average number of observations per landmark was 2.7.

The latitude/longitude coordinates of all 516 landmarks were estimated. The a priori uncertainties on these were set to a fairly large 0.10 (about 10 km) in both the North-South and East-West directions. The landmark radii from the center of Venus were estimated for 24 landmarks located where no Magellan altimetry was available, and for the 107 landmarks seen on both cycles, since the different incidence angles of the measurements from the two cycles made the radii observable. Radii values interpolated from Magellan altimetry we're used as a priori values for the intercycle landmarks, and as fixed values for the rest of the landmarks. The a priori uncertainties on the radii were set at 500 m for landmarks with no altimetry and 100 m for the intercycle landmarks (both 1-sigma). The total number of landmark coordinates estimated was 1163. The Venus pole position and rotation rate were fixed at the current best estimates.

A total of 4442 ground-based Doppler measurements and 629 differenced Doppler measurements were used in the solution. The amount of ground-based tracking data in the cycle 3 arcs was half that in the cycle 1 arcs, however, because for thermal reasons the spacecraft had to be turned away from Earth-pointing for large portions of the orbit (to

hide the spacecraft bus in the shadow of the main antenna). The majority of the available tracking data (70°/0) was at X-band, so it was mostly of good quality. in order to favor the landmark measurements, the ground-based data was deweighted by a factor of 10 from its theoretical accuracy: the X-band measurement uncertainty was set at 0.055 Hz.

The Venus gravity field model used for these solutions was much more precise than that used for the early solutions discussed in Ref. 1. The model used was the 50× 50 field PMGN50F based on all Pioneer Venus Orbiter tracking data and much of Magellan's cycle 4 data[7]. Whereas with earlier gravity models it was necessary to estimate low degree and order gravity coefficients to obtain good fits to the data, this is no longer necessary. Note that this gravity field is also much more precise than that used in the original gg-ound-based ephemeris solutions.

Two other differences from earlier ephemeris improvement solutions concern additional dynamical parameters which were estimated. First, the atmospheric density at periapsis was estimated on every orbit to account for variability in the atmosphere. A priori values for these were set to 3×10^{16} gm/cc. Secondly, three components of velocity change were estimated at every momentum wheel desaturation maneuver, to account for mismodeling of these maneuvers. The a priori values for these were the size of the maneuvers as estimated from telemetry. The total number of parameters estimated, including landmark parameters, dynamic parameters, and the initial states from the 7 arcs was 1316.

Figures 3-5 show residuals for the middle of the three cycle 1 arcs. They arc typical of the other arcs, although this arc had the most landmark measurements (357). Figures 3 and 4 show the landmark range and Doppler residuals, Since the SAR Doppler residuals are normalized by their respective burst durations, which are designed to provide 120 m azimuth resolution, the normalized Doppler scale can be interpreted as along-track landmark position residuals in units of 120 m. The landmark residuals arc quite flat, with no evident trends. The root-mean-squares of the landmark residuals were 24 m in range and 0.37 in normalized Doppler, which are typical of the other arcs as well. The X-band ground-based Doppler residuals, shown in Figure 5, are reasonably flat, but some signatures remain, The plot of the S-band residuals, not shown, is similar, but noisier.

To test the quality of this solution, the sizes of the discontinuities at inter-arc boundaries were measured. The solution ephemeris for each arc was numerically integrated one orbit into the next arc and differenced with the ephemeris from the second arc. Table 1 summarizes these inter-arc differences as measured at a latitude of 65° North, which is near Maxwell; differences at other latitudes are comparable. The first two arc boundaries arc in cycle 1; the remaining are in cycle 3. It may be concluded from this test that the improved ephemeris is very smooth, with relative errors generally less than 80 m, smaller than the radar resolution. The test cannot measure intercycle relative errors, but they must be of the same order of magnitude.

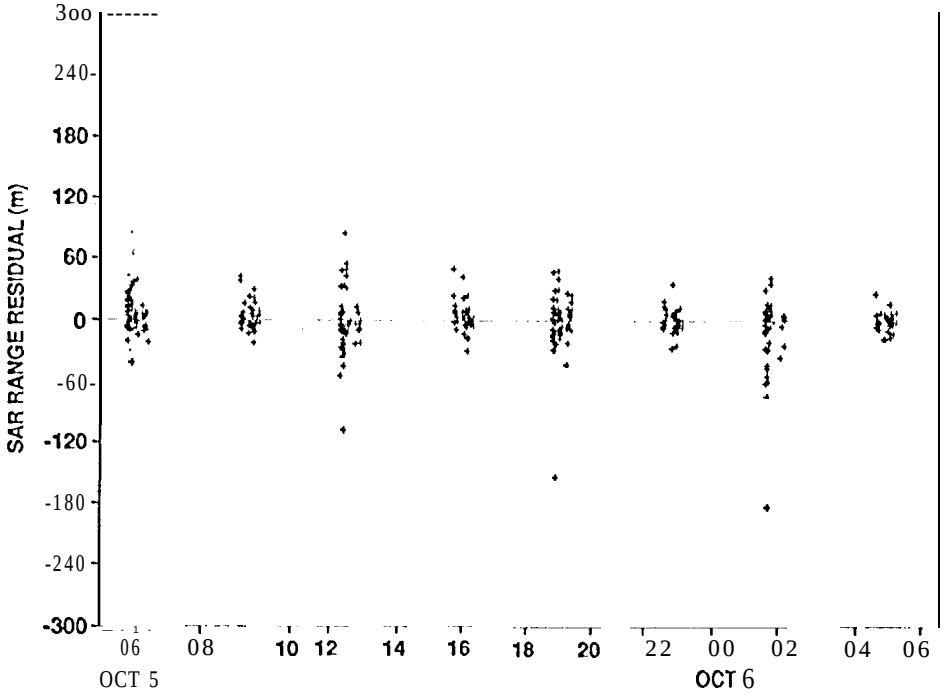

Figure 3. Landmark Range Residuals vs. Time

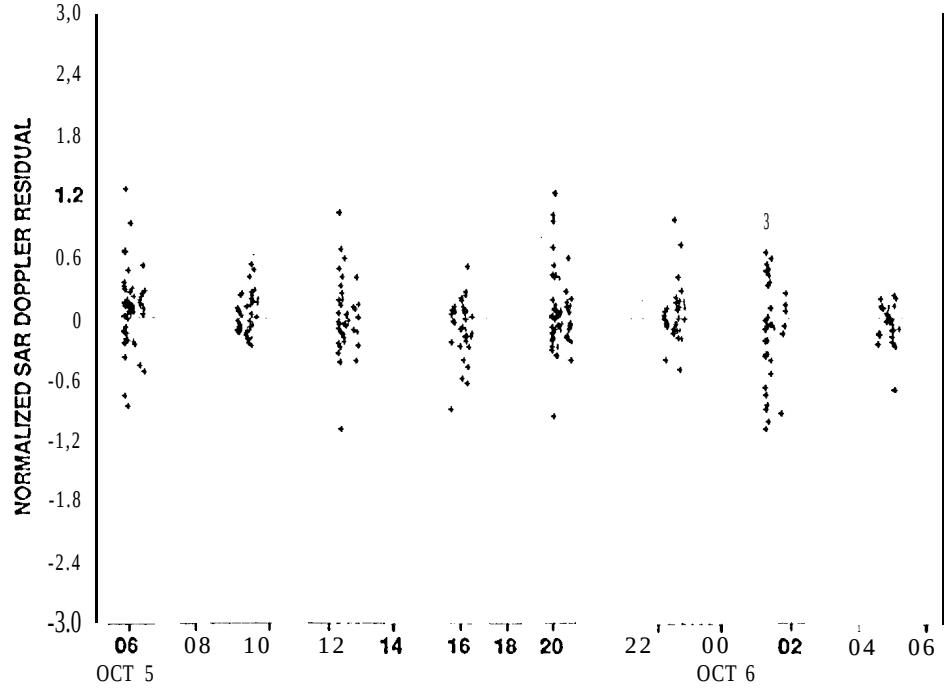

Figure 4. Landmark Normalized Doppler Residuals vs. Time

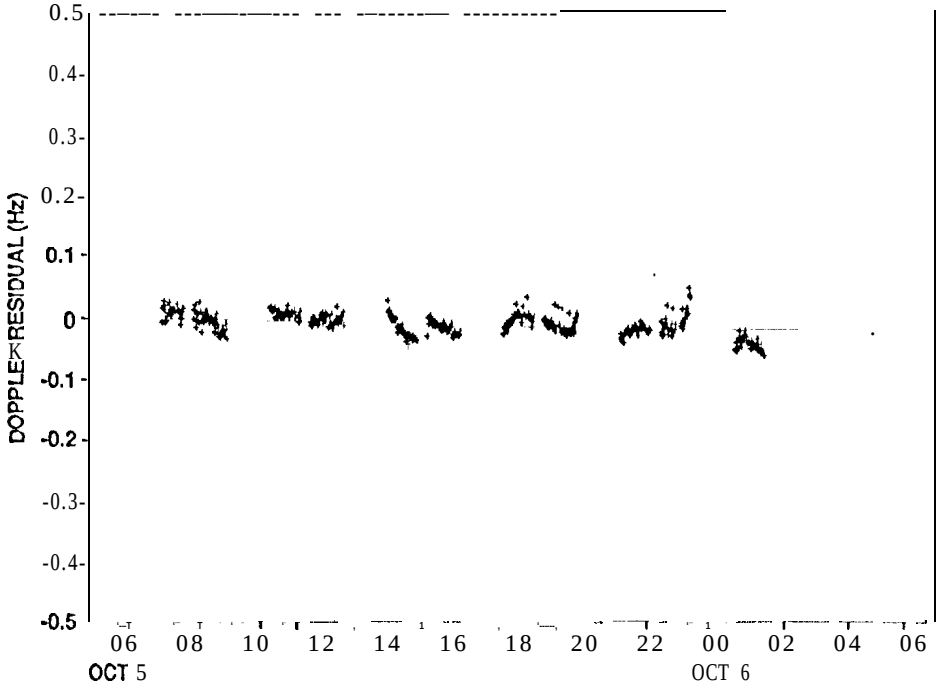

Figure 5. Ground-Based Doppler Residuals (X-Band) vs. Time

Table 1
Measured Discontinuities Across Arc Boundaries at 6S° North

Arcs	Along-Track (m)	Cross-Track (m)	Radial (m)	RSS (m)
1-2	77	15	19	81
2-3	20	12	2	24
4-5	10	37	15	41
5-6	26	17	6	31
6-7	56	118	3	131

The largest discontinuity in the table is 118 m in cross-track between the last two arcs. The reason this is so much larger than the other cross-track discontinuities may be that the last arc had fewer landmark measurements because it contained 2 orbits in which SAR data were deleted duc to the on-board clock calibration errors. Other orbits from this arc might suffer similar clock errors to a lesser extent. The most likely explanation, however, is that the elevations of landmarks in the Maxwell region, as determined from altimetry, were inaccurate in this region, because it is so rugged, and the last arc contained the most mountainous terrain. As further evidence of this explanation, the cross-track discontinuity across the last pair of arcs was a much smaller 40 m in the Gula region at 22° North, where the terrain was smoother. A better ephemeris solution would probably result if none of the elevations of landmarks near Maxwell were fixed at altimetry-derived values. Finally, this inter-arc test may overestimate the size of discontinuities because it ignores errors due to dynamic mismodeling in extending the ephemeris an extra orbit.

For comparison, a similar test of inter-arc boundaries was performed on a set of ephemeris solutions produced using no landmark measurements and a more recent and more precise gravity field (PMGN60D)[7]. Without landmarks, the solutions over the arcs were essentially independent of each other, but were nevertheless very smooth across cycle 1 arc boundaries, with discontinuities even smaller than those in '1'able 1. 1 lowever, the cycle 3 inter-arc boundaries had discontinuities as large as 1 km, probably due to the fact that there was much less ground-based tracking data on the cycle 3 arcs. The absolute differences between the landmark and non-landmark solutions was about 1 km on both cycles, which is probably the level of accuracy of the solution without landmarks.

The improved ephemeris solution using landmarks was used in the reprocessing of the SAR data and regeneration of image swaths for the orbits which covered the two stereo test regions. A simple test of the intercycle relative errors was performed by again measuring landmarks, this time on the new images. The latitude differences between cycles were found to be zero mean. A direct comparison of longitude differences was not possible because of the additional parallax displacement associated with the topography. 1 lowever, a new DEM of the Maxwell region was produced using the new ephemeris and the radar stereo equations approach. The resulting DEM contained no apparent artifacts duc to ephemeris errors. A more quantitative test comparing height determinations in the overlap regions of neighboring orbits led to the conclusion that the cross-track relative ephemeris errors of the improved ephemeris were no larger than about 100 m.

CONCLUSIONS

A large amount of stereo imagery of Venus has been acquired by the Magellan spacecraft. The determination of stereo height information and generation of digital elevation maps is complicated by the fact that stereo processing of spacecraft radar data is very sensitive to ephemeris errors. The Magellan ephemeris determined using only Earth-based tracking data is not precise enough to produce accurate, artifact-free stereo products. SAR landmark measurements can be used to augment Magellan orbit determination to obtain a significantly more precise ephemeris solution, When applied to

multiple data arcs simultaneously, common landmarks tie together otherwise independent pieces of the ephemeris rind reduce relative ephemeris errors bet wcen non-cent iguous data arcs. The technique was applied to overlapping blocks of orbits for a test of Magellan stereo processing. The relative ephemeris errors in the resulting solution were generally smaller than 100 m, a level of precision an order of magnitude better than that of the best solutions using only Earth-based data,

ACKNOWLEDGMENTS

The authors would like to thank M.E. Davies and T.R. Colvin of the RAND Corp., P.G. Ford of MIT, and A.S. Konopliv, W.],. Sjogren, T.C. Wang, G.M. Goell, L. Ploplys, and M. McAuley of JPL for valuable assistance. The research described in this paper was carried out by the Jet Propulsion Laboratory, California institute of Technology, under contract with the National Aeronautics and Space Administration.

REFERENCES

1! Chodas, P. W., T-C. Wang, W.L. Sjogren, and J.E. Ekelund, "Magellan Ephemeris Improvement Using Synthetic Aperture Radar Landmark Measurements," A AS 91-391, AA S/AIAA Astrodynamics Specialist Conf., Durango, Colorado, Aug. 1991

2. Leberl, F. W., *Radargrammetric Image Processing*, Artech 1 louse, Norwood, Massachusetts, 1990

3. Lewicki, S., M. Lee, P. Chodas, and E. DeJong, "Stereo Processing of Magellan SAR Imagery Performed on a Transput er Architecture," 1993 lnt. Geoscience and Remote Sensing Symposium, Tokyo, Japan, Aug. 19934. Engelhardt, D.B., J.B. McNamee, S.K. Wong, F.G. Bonneau, E.J. Graat, R.J. Haw, G.R. Kronschnabl, and M.S. Ryne, "Determination and Prediction of Magellan's Orbit," AAS 91-180, AAS/AIAA Spaceflight Mechanics Meeting, Houston, Texas, Feb. 1991

5. Wang, T. C., J.B. Collier, J.E. Ekelund, and P.J. Breckheimer, "Applications of Square-Root information Filtering and Smoothing in Spacecraft Orbit Determination," Proceedings of the 27th IEEE Conf. on Decision and Control, Austin, Texas, Dec. 1988

6. Bierman, G. J., *Factorization Methods for Discrete Sequential Estimation,* Academic Press, New York, Ncw York, 1977

7. Konopliv, A, S., N.J. Borderies, P.W. Chodas, E.J. Christensen, W.L. Sjogren, B.G. Williams, G. Balmino, and J.P. Barriot, "Venus Gravity and Topography: 60th Degree and Order Model, " *Geophys. Res. Letters*, in press